CDMTCS
Research
Report
Series

Compound Constructions of Minimal Broadcast Networks

Michael J. Dinneen
Department of Computer Science,
University of Auckland, New Zealand

Jose A. Ventura
Department of Industrial and
Manufacturing Engineering,
Pennsylvania State University, USA

Mark C. Wilson
Department of Mathematics,
University of Auckland, New Zealand

Golbon Zakeri
Department of Engineering Science,
University of Auckland, New Zealand

CDMTCS-026
January 1997
Centre for Discrete Mathematics and
Theoretical Computer Science

COMPOUND CONSTRUCTIONS OF MINIMAL BROADCAST NETWORKS

MICHAEL J. DINNEEN, JOSE A. VENTURA, MARK C. WILSON, AND GOLBON ZAKERI

ABSTRACT. Compound methods have been shown to be very effective in the construction of minimal broadcast networks (mbns). Compound methods generate a large mbn by combining multiple copies of an mbn G using the structure of another mbn H. Node deletion is also allowed in some of these methods. The subset of connecting nodes of G has been defined as solid h-cover by Bermond, Fraigniaud and Peters, and center node set by Weng and Ventura. This article shows that the two concepts are equivalent. We also provide new properties for center node sets, including bounds on the minimum size of a center node set, show how to reduce the number of center nodes of an mbn generated by a compound method, and propose an iterative compounding algorithm that generates the sparsest known mbns in many cases.

1. INTRODUCTION

Communication in networks is a process whereby a set of messages, generated by a set of originators, is transferred to a set of receivers. The nodes of the network are the possible originators and receivers of the messages, and the edges are the communication lines which allow the direct transmission of messages between certain pairs of nodes. There is a wide range of network design problems in communication networks which are differentiated by placing constraints upon the set of messages, the originators, the receivers, the edges, the network topology, and the transmission characteristics of the network [1, 9, 12, 13, 18, 21]. A communication network can be modeled as a connected graph $G = (V(G), E(G))$ without loops or parallel edges, consisting of a set of nodes $V(G)$ with cardinality $v(G)$, and a set of undirected edges $E(G)$ with cardinality $e(G)$.

Broadcasting is a special type of network communication in which a single message, originated at any node, is transmitted to all the other nodes of the network. Broadcasting is usually required to be completed as rapidly as possible by a sequence of transmissions through the communication lines. It is assumed that broadcasting is carried out under the following three constraints [7, 8]: (i) each transmission requires one unit of time, (ii) a node can make at most one transmission in one time unit, and (iii) a node can only transmit the message to its neighbors (two nodes are called neighbors if they are connected by an edge). Thus, in one time step, the number of informed nodes can at most be doubled. This implies that after m time steps the number of nodes that have received the message, including the originator, is at most 2^m. The *broadcast time* $b(G)$ of a graph G is the minimum number of time steps in which broadcasting can be achieved in G regardless of the originator of the message. From the above it is clear that $b(G) \geq \lceil \log_2 v(G) \rceil$. A *minimal broadcast network* (mbn) is defined to be a communication network G satisfying $b(G) = \lceil \log_2 v(G) \rceil$. Complete graphs are obviously mbn's.

Suppose that a node u in a network G is the originator of the message. A *broadcast protocol* (or broadcast tree) $P(G, u)$ is a rooted spanning tree in which the originator u is the root and all the nodes are labeled by their receiving times. In a broadcast protocol, each edge is used exactly once and the message is always transmitted from parent to child. In order that a network G be an mbn,

Date: January 7, 1997.
1991 *Mathematics Subject Classification.* Primary 90B12. Secondary 05C90, 68R10, 94A05.
Key words and phrases. broadcasting, communication network, minimal broadcast graph, graph compound.

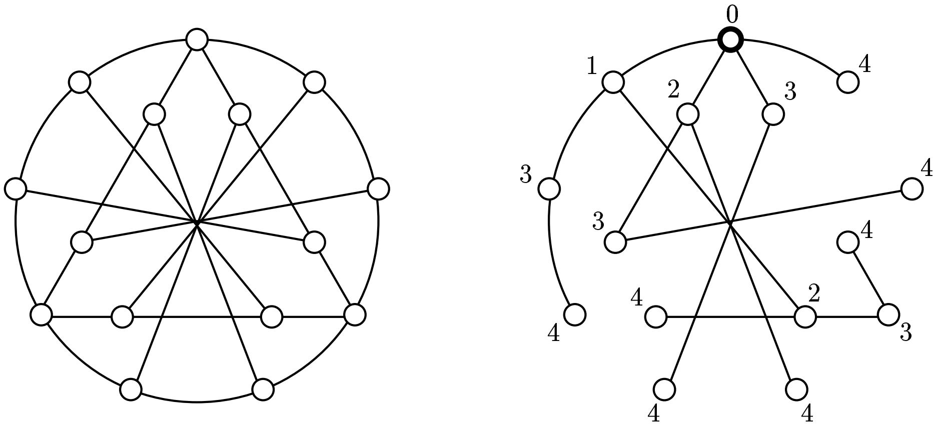

(a) Obn with 15 nodes. (b) Broadcast protocol.

FIGURE 1. An obn with 15 nodes and a selected protocol.

each node in the network must have a broadcast protocol that can be completed in $\lceil \log_2 v(G) \rceil$ time steps.

The problem of recognizing whether an arbitrary network is an mbn is $\mathcal{NP}$-complete [7]. An *optimal broadcast network* (obn) is an mbn with the minimum possible number of edges for its given number of nodes, and the *broadcast function* $B(n)$ is defined to be the number of edges of every obn with n nodes. There is no known feasible method for determining $B(n)$ for an arbitrary value of n. Farley et al. [8] showed that hypercubes are obn's and so $B(2^m) = m2^{m-1}$ for $m \geq 0$. Khachatrian and Harutounian [14] and Dinneen et al. [6] proved independently that $B(2^m-2) = (m-1)(2^{m-1}-1)$ for $m \geq 2$. Farley et al. also determined the values of $B(n)$, for $1 \leq n \leq 15$. Bermond et al. [2] and Weng and Ventura [23] published known values of $B(n)$ for $17 \leq n \leq 63$. Recently, Sacle [20] gave lower bounds on $B(2^m - k)$, for $m \geq 3$ and $3 \leq k \leq 6$, Figure 1 shows an example of an obn with 15 nodes and 24 edges and one of its broadcast protocols.

Direct construction of mbn's is a difficult process, since in the worst case one must check that every node has a broadcast protocol taking time $\lceil \log_2 v(G) \rceil$. Thus most authors have concentrated on constructions which combine several known mbn's to create new ones.

In 1979 Farley [7] introduced mbn's and obn's, proposed a recursive algorithm to construct mbn's with an arbitrary number of nodes n, and showed that the number of edges of the mbn's produced by his algorithm is bounded by $(n/2)\lceil \log_2 n \rceil$. Chau and Liestman [4] developed an algorithm which constructs mbn's by interconnecting 5, 6 and 7 smaller mbn's. They also improved Farley's bound on $B(n)$, for n in the range $(2^{m-1}, 7 \cdot 2^{m-3})$, $m \geq 1$. Gargano and Vaccaro [10] proposed three algorithms based on the interconnection of hypercubes of small dimension to build up larger mbn's. Chen [5] presented a method similar to the second algorithm of Gargano and Vaccaro, and then suggested the recursive application of his first method to construct larger mbn's. In Grigni and Peleg's algorithm [11], hypercubes and generalized Fibonacci numbers are used to construct mbn's with $O(L(n)n)$ edges, where $L(n)$ is the number of leading 1's in the binary representation of $n-1$. In practice, however, the other methods seem to require fewer edges. Bermond et al. [3] proposed four methods to construct mbn's for $18 \leq n \leq 63$. Ventura and Weng [22] developed a method based on the concepts of aggregated nodes and aggregated edges (which are used to replace ordinary nodes and edges, respectively, of known obn's, for $9 \leq n \leq 15$) to construct sparse mbn's. The central idea of all the methods discussed so far is to produce larger mbn's by combining small known obn's or mbn's using as few edges as possible without violating the constraint of being able to broadcast in minimum time from any node.

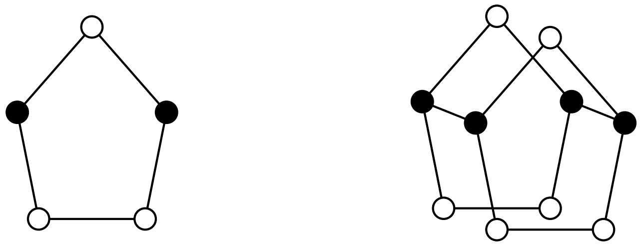

(a) Obn with 5 nodes. (b) Obn with 10 nodes.

FIGURE 2. Constructing an obn with 10 nodes by compounding.

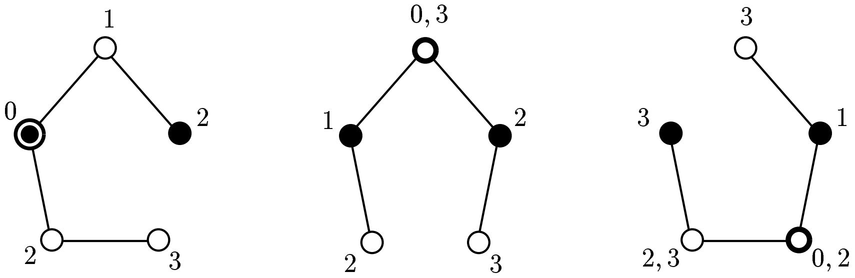

FIGURE 3. The official broadcast protocols for the obn in Figure 2(a).

More recently another class of combination methods using compound graphs has been developed by Bermond, Fraigniaud and Peters [2], generalizing a construction of Khachatrian and Harutounian [14]. A more general method, which allows for systematic vertex deletion, was proposed by Weng and Ventura [23]. A key ingredient of this last method, called the *doubling procedure*, is a center node set, defined via so-called official broadcasting.

In this article, we investigate the constructions of [2] and [23], treating official broadcasting and center node sets in more detail. Iterative algorithms based on these constructions are presented and analyzed. Computational results are obtained which improve most of the best known bounds on numbers of edges and size of center node sets achieved by previous methods.

In Section 2 we recall the definitions of official broadcasting and center node set and demonstrate the equivalence of the latter concept and that of solid h-cover. We derive bounds on the minimal size of a center node set. Section 3 focuses on ways of reducing the center node sets generated by the compounding procedures. The short Section 4 discusses the general framework of the iterative algorithms and shows the limitations of the center node reductions. Section 5 has three subsections. In the first of these we discuss our practical implementation of the iterative algorithms. In the second part the initial input used by the algorithms is verified. An important feature here is an efficient and accurate calculation of good center node sets for known mbn's. In the third part, we present computational results that compare our refined algorithm with previously described methods. A table containing the known values of $B(n)$ and the best upper bounds on $B(n)$, for $17 \leq n \leq 127$, is also presented. Finally, some open problems and directions for future research are discussed in Section 6.

2. Preliminaries

This section lays out the basic notation and definitions concerning center node sets which will be used in the rest of the paper. The idea of official broadcasting, described below, leads naturally to the definition of a center node set. We establish the equivalence of the center node sets of Weng and Ventura [23] and the solid h-covers of Bermond et al. [2]. In so doing we recall the compounding methods introduced in these papers. We obtain elementary bounds for the minimum size of center node sets which will prove useful later.

The following definition will be used throughout. Let (n, k, i) be a triple of integers with $n > 0$ and $0 \leq i < k$. Then

$$\lceil \log_2(nk - i) \rceil \leq \lceil \log_2(nk) \rceil \leq \lceil \log_2 n \rceil + \lceil \log_2 k \rceil.$$

We say that the triple (n, k, i) *satisfies the broadcast condition* if equality holds in both inequalities. If $i = 0$ we say simply that that (n, k) satisfies the broadcast condition.

2.1. **Official broadcasting and center node sets.** The *compound* of a graph G into another graph H relative to a set $S \subseteq V(G)$, denoted by $G_S[H]$, is the graph obtained by replacing each node of H with a copy of G, and each edge of H by a matching between the corresponding copies of S. Thus, $V(G_S[H]) = V(G) \times V(H)$ in a natural way.

In [14] a compounding algorithm was developed in which G is an mbn restricted to have a maximum degree bounded by $\lceil \log_2 v(G) \rceil - 1$, S is a vertex cover of G and $H = K_2$.

In [2] a more general compounding procedure was presented, in which G and H are general mbn's satisfying the broadcast condition, and S was what was called a solid h-cover of G (see subsection 2.2 below). Figure 2 illustrates this method by generating an obn with 10 nodes from two copies of an obn with 5 nodes. In this example S is a solid 2-cover defined by the black nodes in Figure 2(a) and $H = K_2$.

It is not possible to generate mbn's with n nodes for every n by this procedure (for example, when n is prime). Weng and Ventura [23] proposed a generalized compounding algorithm, which they called the *doubling procedure*, which includes the method of [2] as a special case. The doubling procedure's extra generality mainly arises from the fact that nodes may be deleted in certain copies of G. The method constructs an mbn from given mbn's G and H, and an integer i with $0 \leq i \leq v(H) - 1$, such that $(v(G), v(H), i)$ satisfies the broadcast condition. The subset S of nodes of G used for connecting the copies of G was called a center node set in [23].

Weng and Ventura introduced the concepts of official broadcasting and center node sets in order to describe the doubling procedure. In *official broadcasting*, certain nodes of the network, called *center nodes*, are given the authority to make a message official. A message, originated at any node, must be made official during the broadcast protocol. The official message must then be transmitted to all the nodes of the network. During the broadcast protocol a message is official if it has been officialized by a center node; otherwise, it is unofficial. It is assumed that an unofficial message becomes official immediately after it arrives at a center node. In official broadcasting, where all the nodes must receive an official message, a center node will only receive one message, so that if the incoming message is unofficial, it will be officialized immediately after its arrival. A non-center node may receive one or two messages. In the first case, the message must be official. In the second case, the first message must be unofficial and the second one official. In addition, in official broadcasting, it is possible for a non-center node to send an unofficial message to a neighbor and receive an official message during the same time step.

A more formal description is as follows. An *official broadcast protocol* for a node u with respect to a set $S \subseteq V(G)$ in a graph G, denoted by $P(G, u; S)$, is a spanning subgraph of G containing all of its nodes, in which the nodes are labeled by one or two receiving times, all of which are at most $b(G)$. If a node has two receiving times, it must not belong to S, the first receiving time is for the unofficial message and the second one for the official message. If $u \in S$, each node is labeled by a single receiving time, and the official broadcast protocol is a spanning tree rooted at u, that is, just an ordinary broadcast protocol.

Given an official protocol $P = P(G, u; S)$ we define the *unofficial part* P_u of P to be the tree rooted at u induced by all edges that transmit an unofficial message and the corresponding nodes.

The *official part* P_o of P is a forest of rooted trees induced by all edges which transmit an official message. These trees are rooted at nodes in S which receive an unofficial message. Denote the set of all such nodes by $V_{cu}(P)$. The forest P_o spans G.

In particular, if every node of a graph G has an official protocol of time $b(G)$ with respect to S, then S is called a (minimal) *center node set* of G. Clearly if S is a center node set and $S \subseteq T \subseteq V(G)$ then T is also a center node set. In this paper we shall be interested only in the case where G is an mbn. In this case, unless otherwise stated all protocols are assumed to take $\lceil \log_2 v(G) \rceil$ time steps.

If P is an ordinary broadcast protocol, then a node is *idle at time* t if it has received the message at time $t-1$ and does not transmit the message at time t. It is important to note that a node u can only be a non-center node if in some broadcast protocol for u, some node is idle and can send the official message back to u. In practice we construct official protocols by using ordinary protocols and showing that there are enough idle nodes to inform all non-center nodes that receive an unofficial message.

The notion of official broadcasting was introduced purely as a way to describe the set of connecting nodes in a compound mbn. However it may have other applications, such as minimizing costs in the design of networks in which message authentication is required. Hence obtaining the smallest possible center node sets of a given graph is a problem of great interest.

Given a graph G, define the *center node number* $cn(G)$ to be the minimum size of all center node sets for G. A center node set of size $cn(G)$ is called an *optimal center node set (ocns)* of G.

There may exist multiple ocns's for a given graph G, and the cardinality of ocns's of non-isomorphic graphs with n nodes and m edges may be different. There is no known polynomial-time algorithm for computing $cn(G)$ for an arbitrary graph G. In the obn with 5 nodes presented in Figure 2(a), the black nodes are center nodes and the white nodes are non-center nodes. Actually, the two center nodes define an ocns for the obn. Figure 3 shows the three different official broadcast protocols for this obn with 5 nodes.

2.2. **Compound methods.** We can now describe the *doubling procedure* of [23]. Fix an mbn G with center node set S. If $S \neq V(G)$ then choose a non-center node v of G (usually of minimal degree). Otherwise let v be any (center) node of S. Construct a new network G^v by deleting v and all its incident edges from G, and adding the required edges to form a clique among the neighbors of v. Let $S^v = S \setminus \{v\}$.

Let H be an mbn. For a fixed integer i with $0 \leq i \leq v(H) - 1$, we can construct a network $\mathcal{G}$ by connecting $v(H) - i$ copies of G and i copies of G^v as follows. For each fixed $s \in S^v$, connect all $v(H)$ copies of s to form a copy H_s of H. If v is a center node of G, the $v(H) - i$ copies of v are connected to form an mbn H^* with $v(H) - i$ nodes. If v is a non-center node then let H^* be the empty graph. The procedure is illustrated in Figure 4. We shall at times refer to *vertical* (within G or G^v) or *horizontal* (within H or H^*) broadcasting; these terms should be interpreted as illustrated in that figure.

Theorem 2.1 (The Doubling Procedure [23]). *Let G and H be mbn's and let S be a center node set for G. Suppose that $(v(G), v(H), i)$ satisfies the broadcast condition.*

(i) *The graph $\mathcal{G}$ constructed by the doubling procedure is an mbn.*
(ii) *The set $\bigcup_{s \in S^v} v(H_s) \cup v(H^*)$ is a center node set for $\mathcal{G}$.*

□

We point out here that the example on page 292 of [23] is invalid, since $(5, 17)$ does not satisfy the broadcast condition.

The doubling procedure has the following important special case.

Theorem 2.2 (The Ordinary Compounding Method (cf. [2])). *Let G and H be mbn's and let S be a center node set for G. Suppose that $(v(G), v(H))$ satisfies the broadcast condition.*

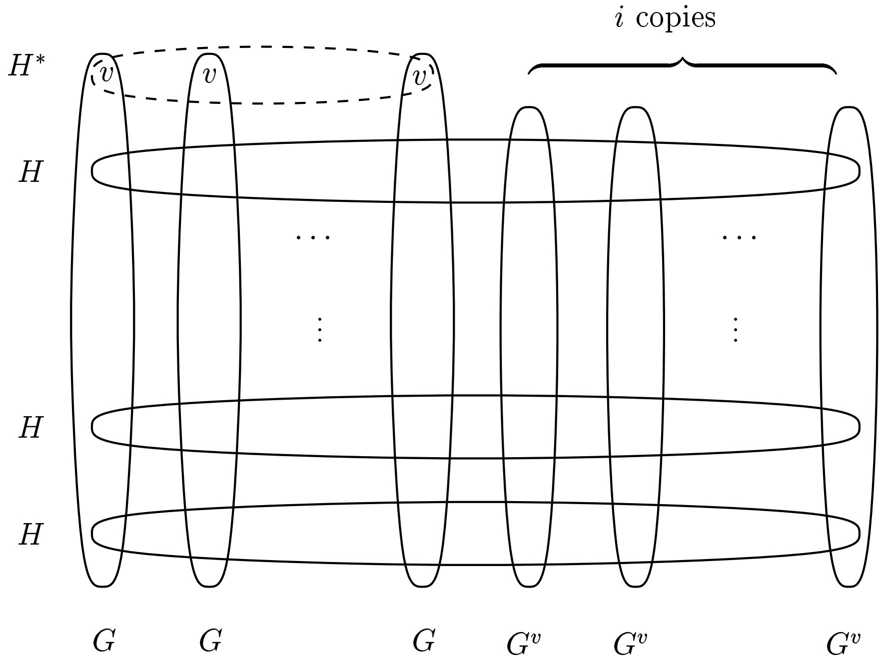

FIGURE 4. The doubling procedure of Weng and Ventura (see [23]).

(i) *The compound graph $\mathcal{G} = G_S[H]$ is an mbn.*
(ii) *The product $S \times V(H)$ is a center node set for $\mathcal{G}$.*

□

In [2] similar results to Theorem 2.2(i) are presented. The only difference is that S is required to be a solid 1-cover or solid 2-cover as defined in that paper. We now establish the connection between the two papers.

Recall from [2] that for a given mbn G, a subset S of nodes is a *solid 1-cover* if S is a vertex cover of G and for every $u \notin S$, there is a broadcast protocol for u such that at least one neighbor of u is idle at some time during the broadcast.

Proposition 2.3. *A subset S of $V(G)$ is a solid 1-cover if and only if S is both a vertex cover and a center node set.*

PROOF. Let S be a solid 1-cover for G. If $u \in S$ then there is an official protocol for u by definition. If $u \notin S$ then every neighbor of u belongs to S. Thus in the protocol guaranteed by the definition of solid 1-cover, for every node v except u, each message is official on arrival at v. It remains only to inform u officially, and this is done by the idle neighbor supplied by the definition.

For the reverse direction, if S is a center node set and vertex cover then for each $u \notin S$ there is an official protocol for u with respect to S. The node which sends the official message to u is necessarily in S and must be idle in some broadcast protocol for u. □

In [2] it is stated that one can define solid h-covers for every $h \geq 0$ and the conclusion of Theorem 2.2(i) will hold with S assumed to be a solid h-cover. An inspection of the proof shows that it is necessary to *define* a solid h-cover to be a center node set which is an h-cover, that is, a path cover for all paths of length h (note that a vertex cover is a 1-cover). Since every center node set is an h-cover for some h, the two concepts then coincide.

It is possible that one might obtain an mbn by compounding with respect to a set of nodes that is not a center node set (this would require a different method of proof to that in [23] and [2]). We believe that this is not the case and that center node sets are crucial for such a result.

Conjecture 2.4. Let G and H be mbn's and let $S \subseteq V(G)$. If $G_S[H]$ is an mbn then S is a center node set for G.

Even when $H = K_2$ the above conjecture seems rather difficult.

2.3. **General bounds on center node sets.** We now present some elementary results which are useful in obtaining bounds on $cn(G)$.

Perhaps the easiest way to obtain a good center node set is the following widely applicable result.

Proposition 2.5. *Let G be an mbn of broadcast time $m = \lceil \log_2 v(G) \rceil$ and let $S \subseteq V(G)$. If S is a vertex cover for G such that every neighbor of every node not in S has degree at most $m-1$, then S is a center node set for G.*

PROOF. Let v be a node not in S. In every broadcast protocol P for v the first neighbor receiving the unofficial message must be idle at time no later than m and can therefore send the official message back to v. All other nodes already have an official message and so this yields an official protocol for v. □

We now proceed to derive a lower bound for $cn(G)$ that is valid under special conditions on $v(G)$. Observations similar to the following have been made by many authors (see for example [11, 16]). This result holds in the case when v is never idle, i.e. $l = 0$.

Proposition 2.6. *Let G be an mbn with broadcast time m. Let v be a node of G. Suppose that in a broadcast protocol for v, v itself is idle at times $t_1 < t_2 < \cdots < t_l$.*

(i) $v(G) \leq 2^m + 1 - 2^l$.

(ii) *Equality occurs if and only if v is idle for the last l time steps, and no other node is idle.*

PROOF. Denote by $I(i)$ the number of nodes informed by time i. Then $I(t_1 - 1) \leq 2^{t_1 - 1}$. Since v is idle at time t_1 and the size of the forest formed by the informed nodes other than v can at most double at time t_1, we have $I(t_1) \leq 2(2^{t_1-1}) + 1 = 2^{t_1}(1 - 2^{-t_1})$. An easy induction yields $I(t_l) \leq 2^{t_l}(1 - (2^{-t_1} + \cdots + 2^{-t_l})$. Thus $I(m) \leq 2^{m-t_l} I(t_l) \leq 2^m(1 - (2^{-t_1} + 2^{-t_2} + \cdots + 2^{-t_l}))$. This last expression is clearly maximized when and only when $t_1 = m + 1 - l, \ldots, t_m = m$ and the corresponding maximal value is $2^m + 1 - 2^l$. □

Let $\delta(G)$ denote the minimum degree of a node of G.

Corollary 2.7. *Let G be an mbn with broadcast time m. Then*

$$\delta(G) \geq m - \lfloor \log_2(2^m + 1 - v(G)) \rfloor.$$

PROOF. If an originating node of G has degree $k \leq m$ then it must be idle at least $m - k$ times. Thus $v(G) \leq 2^m + 1 - 2^{m-k}$ by Proposition 2.6. Solving for k yields $k \geq m - \log_2(2^m + 1 - v(G))$ and the result follows since $\delta(G) \in \mathbb{Z}$. □

A weaker bound, with the floors replaced by ceilings is contained in Theorem 2.2 of [11]. The latter bound has a simple interpretation in terms of the binary expansion of $v(G) - 1$, but it seems that the bound above does not.

The corollary yields the important special cases that mbn's on 2^m, $2^m - 1$ and $2^m - 2$ nodes have minimum possible degrees m, $m - 1$ and $m - 1$, respectively.

For each $n \in \mathbb{Z}$ let $V_n(G)$ denote the set of nodes of G of degree n.

Proposition 2.8. *Let G be an mbn with $2^m + 1 - 2^{m-\delta}$ nodes for some δ with $0 \leq \delta \leq m$.*

(i) *Let v be a node of degree δ. Then v is contained in every center node set of G. Hence $cn(G) \geq |V_\delta(G)|$.*

(ii) *If $\delta < m$ and V_δ is a vertex cover then $cn(G) = |V_\delta(G)|$.*

PROOF. By Proposition 2.6(ii), in an official broadcast protocol for v with respect to any subset S of $V(G)$, the only node that can possibly be idle is v itself (for the last $m-\delta$ time steps). Thus no node can send an official message back to v and so necessarily $v \in S$. This proves (i), and (ii) follows directly from Proposition 2.5. □

Taking $\delta = m$ and $\delta = m-1$ respectively we see that (i) above generalizes Theorems 2.1 and 2.3 of [23].

Finally, general bounds on $cn(G)$ are given by the next result. The lower bound is attained by all complete graphs. We believe that the degree hypothesis in (iii) is redundant, but do not have a proof of this.

Proposition 2.9. *Let G be an mbn with broadcast time m.*

(i) *If $v(G) = 2^m$ then $cn(G) = v(G)$, whereas if $v(G) < 2^m$ then*
$$m - \lfloor \log_2(2^m - v(G)) \rfloor \leq cn(G).$$
(ii) *If equality occurs in (i) then G has an ocns which is a vertex cover.*
(iii) *If there is a subset R of nonadjacent nodes of $V(G)$ all of whose neighbors have degree at most $m-1$, then*
$$cn(G) \leq v(G) - |R|.$$

PROOF. If $v(G) = 2^m$ then no node can be idle during a broadcast so all nodes are center nodes. If $v(G) < 2^m$ then suppose that k is the size of a center node set for G. If $k > m$ the result follows trivially so we may assume $k \leq m$. Let v be a non-center node. Then the maximum number of nodes with the official message at time k is $2^k - 1$, and this occurs only if all neighbors of v are center nodes. Thus the maximum size of the official forest at time m is $2^{m-k}(2^k - 1)$. This implies that $2^m - 2^{m-k} \geq v(G)$. Solving for k yields $k \geq m - \log_2(2^m - v(G))$ and (i) follows since $k \in \mathbb{Z}$. Part (ii) is immediate since all neighbors of v are center nodes. Part (iii) follows from Proposition 2.5 since $V(G) \setminus R$ is a vertex cover for G. □

Of course for a given $v(G)$, a reduction in the number of edges (while keeping the resulting graph an mbn) may force an increase in the number of center nodes. For example, Conjectures 1 and 2 in [23] say essentially that for each obn G of order $v(G) = 2^m - 1$: (i) all nodes have degree m or $m-1$, (ii) there are $\lceil v(G)/(m+1) \rceil$ nodes of degree m, and (iii) the set of all nodes of degree $m-1$ is a vertex cover of G. Thus in this case $cn(G)$ is asymptotic to $\frac{m}{m+1}(2^m - 1)$, whereas the lower bound above is only $m+1$. This playoff between edges and center nodes is of crucial importance in the compounding methods.

3. Reducing center node sets in compound methods

In this section we consider ways of reducing the center node sets generated by the compounding procedures under consideration. The bounds on $cn(\mathcal{G})$ given by Theorems 2.1 and 2.2 are very poor in general and they deteriorate upon further compounding.

We first treat the simpler special case of Theorem 2.2. The next result reduces the center node sets obtained by applying that theorem.

Theorem 3.1. *Adopt the notation and hypotheses of Theorem 2.2. Let T be a center node set for H. Then $\mathcal{S} = S \times T$ is a center node set for $\mathcal{G}$.*

PROOF. Let u be a node of $\mathcal{G}$. If u belongs to some copy H_s of H then first broadcast officially in H_s with respect to T_s. After $\lceil \log_2 v(H) \rceil$ time steps all nodes in H_s have the official message. Now broadcast vertically in the appropriate copy of G. All nodes of $\mathcal{G}$ are officially informed by time $\lceil \log_2 v(G) \rceil + \lceil \log_2 v(H) \rceil$.

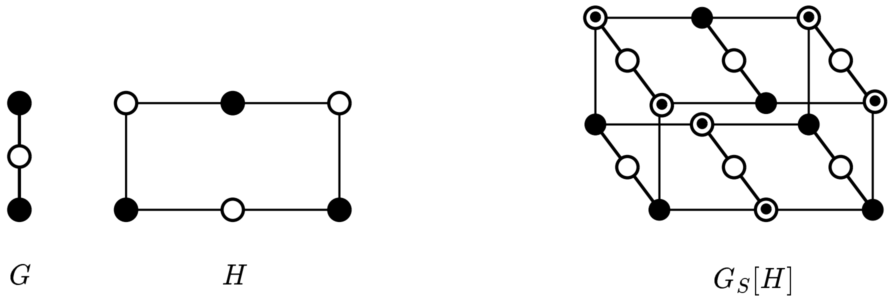

FIGURE 5. Compounding with center node reduction.

Now suppose that u is a node of some copy G_i of G and does not belong to any H_s, i.e. u does not belong to any copy S_i of S. Let P be an official broadcast protocol for u in G with respect to S. Recall the notation $V_{cu}(P)$ from Section 2. Broadcast in G_i according to P_u, ending at $V_{cu}(P)$. On receiving the message each element of $V_{cu}(P)$ then broadcasts horizontally in its corresponding copy of H, and this broadcast is official and takes $\lceil \log_2 v(H) \rceil$ time steps. After this is completed, each node in each copy of $V_{cu}(P)$ continues according to P_o. Since it takes $\lceil \log_2 v(G) \rceil$ time steps to perform official broadcasting in G according to P_u and P_o, the total time needed to broadcast officially in $\mathcal{G}$ is $\lceil \log_2 v(G) \rceil + \lceil \log_2 v(H) \rceil$. □

Figure 5 exemplifies this result. Here S is the set consisting of the two black nodes of G. The circled black nodes in $G_S[H]$ are those center nodes produced by Theorem 2.2 which can be removed according to Theorem 3.1.

We now collect some data about the above construction for later reference.

Proposition 3.2. *The following statements hold for the ordinary compounding method.*

$$\begin{aligned} v(\mathcal{G}) &= v(G)v(H) \\ e(\mathcal{G}) &= v(H)e(G) + |S|e(H) \\ |\mathcal{S}| &= |S||T| \end{aligned}$$

□

We now describe a slight generalization of the doubling procedure of Theorem 2.1. We do not require all H_s to be isomorphic, only that the H_s are mbn's and have $v(H_s)$ all equal to a common integer which we denote $v(H)$. The $e(H_s)$ need not have the same value. Also, it is only necessary to assume the broadcast condition on $(v(G), v(H))$ is satisfied. Finally, H^* need not be an mbn; it is only necessary that broadcast time of H^* is at most that of each H_s. We refer to this construction as the *generalized compounding method.*

Under these weakened hypotheses we can prove the result below. Note that part (ii) drastically reduces the size of the center node set obtained in Theorem 2.1. The condition on U in (ii) is always satisfied if H^* is an mbn and U a center node set for H^*.

Theorem 3.3 (The Generalized Compounding Method)**.** *Suppose that $(v(G), v(H))$ satisfies the broadcast condition.*

(i) *The $\mathcal{G}$ constructed by the generalized compounding method is an mbn.*
(ii) *Let T_s be a center node sets for H_s. Choose $U \subseteq V(H^*)$ so that official broadcast in H^* with respect to U can be completed in $\lceil \log_2 v(H) \rceil$ time steps. Then $\mathcal{S} = \bigcup_{s \in S^v} T_s \cup U$ is a center node set for $\mathcal{G}$.*

PROOF. Recall that it was proved in Theorems 3.1 and 3.2 of [23] that official broadcast in G^v can be carried out with respect to S^v in $\lceil \log_2 v(G) \rceil$ time steps. Furthermore, for each (official) protocol P originating at any node of G other than v, there is a modification P^v of P which works for G^v.

Let u be a node of $\mathcal{G}$. We show that official broadcast with respect to $\mathcal{S}$ is possible from u in $\lceil \log_2 v(G)v(H) \rceil$ time steps.

First suppose that u belongs to some H_s. As in the proof for the special case above, official broadcast within H_s with respect to T_s can be completed in $\lceil \log_2 v(H) \rceil$ time steps. At this stage all nodes in H_s have an official message. In the remaining $\lceil \log_2 v(G) \rceil$ time steps, each node in H_s broadcasts vertically inside its copy of G or G^v. By this time, all nodes in $\mathcal{G}$ are officially informed, proving the result in this case.

Next suppose that u is a node of H^*. Then all nodes of G are center nodes and so all nodes in $\mathcal{G}$ belong to H^* or some H_s. Also u belongs to a unique copy G_i of G. First broadcast vertically in G_i, which takes $\lceil \log_2 v(G) \rceil$ time steps, at the end of which all nodes in G_i have the message. Now broadcast horizontally and officially in the appropriate H_s or H^*. This can be done in $\lceil \log_2 v(H) \rceil$ time steps and so all nodes are officially informed in the correct time.

Suppose now that u is a node of some copy G_i of G but not a node of any H_s. As in the proof of Theorem 3.1, broadcast according to the unofficial part P_u of some official protocol P, ending at $V_{cu}(P)$. Each center node in $V_{cu}(P)$ broadcasts horizontally and officially in its corresponding copy of H. Then each copy t of a node in $V_{cu}(P)$ continues as in P_o (resp. $(P^v)_o$) if t belongs to a copy of G (resp. G^v). As in the special case above this can all be completed in the required time.

Finally, suppose that u is a node of some G^v and not a node of any H_s. Let P be an official protocol for u in G. Broadcast in G^v according to P^v ending in $V_{cu}(P)$. Let Q be the subprotocol followed so far. Let V_1 (respectively V_2) be the subset of $V_{cu}(P)$ consisting of nodes which have (respectively do not have) v as an ancestor. For nodes in V_2 then we have so far followed only $(P^v)_u$, and this is precisely P_u for such nodes. However, for each $t \in V_1$, t may have descendants in $V_{cu}(P)$ which are not in $V_{cu}(P^v)$ and thus Q will include a part of $(P^v)_o$ as far as t and its descendants are concerned.

Now broadcast horizontally and officially in the appropriate H_s, which takes $\lceil \log_2 v(H) \rceil$ time steps. Each copy of each $t \in V_{cu}(P)$ then follows this with broadcasting until the end of P or P^v as appropriate. Now Q takes the same time as $(P^v)_u$ for nodes in V_2. This is the same as the time taken by P_u for such nodes. Also Q takes the same time as P_u for nodes in V_1. Hence Q takes the same time as P_u for all nodes in $V_{cu}(P)$. Since P^v takes $\lceil \log_2 v(G) \rceil$ time steps, and P_u and P_o also use this many time steps, all nodes in $\mathcal{G}$ are informed by the required time. □

We collect some data about the above construction. Here we let $d = \deg v$, and for the given center node set, δ_n and δ_c denote the minimal degree of a non-center and center node respectively.

Proposition 3.4. *The following statements hold for the generalized compounding method.*

$$v(\mathcal{G}) = v(G)v(H) - i$$

$$e(\mathcal{G}) = \begin{cases} v(H)e(G) + id(d-3)/2 + \sum_{s\in S} e(H_s), & \text{if } S \neq V(G) \\ v(H)e(G) + id(d-3)/2 + \sum_{s\in V(G)\setminus\{v\}} e(H_s) + e(H^*), & \text{if } S = V(G) \end{cases}$$

$$|\mathcal{S}| = \begin{cases} \sum_{s\in S} |T_s|, & \text{if } S \neq V(G) \\ \sum_{s\in V(G)\setminus\{v\}} |T_s| + |U|, & \text{if } S = V(G) \end{cases}$$

$$\delta_n(\mathcal{G}) \leq \begin{cases} \min_{s\in S}\{d, \delta_n(H_s) + \delta_c(G)\}, & \text{if } S \neq V(G) \\ \delta_c(G) + \delta_n(H^*), & \text{if } S = V(G) \text{ and } U \neq V(H^*) \\ v(\mathcal{G}) - v(H) + \delta_n(H_s), & \text{if } S = V(G), U = V(H^*) \text{ and } T_s \neq V(H_s) \end{cases}$$

$$\delta_c(\mathcal{G}) \leq \begin{cases} \min_{s\in S}\{\delta_c(G) + \delta_c(H_s)\}, & \text{if } S \neq V(G) \\ \delta_c(G) + \delta_c(H^*), & \text{if } S = V(G) \end{cases}$$

PROOF. Only the entries for δ_n and δ_c require explanation, the first three being exact.

If $S \neq V(G)$ then v is a non-center node of G which remains a non-center node in $\mathcal{G}$. Furthermore v has no new neighbors in $\mathcal{G}$ and so $\delta_n(\mathcal{G}) \leq d$. Now fix $s \in S$ so that $\deg_G s = \delta_c(G)$ and let u be a non-center node of H_s. The copy of s at u is a non-center node of $\mathcal{G}$ of degree at most $\delta_n(H_s) + \delta_c(G)$.

If $S = V(G)$ and $U \neq V(H^*)$ then let u be a node of H^* of degree at most $\delta_n(H^*)$ and let s be a node of G of degree at most $\delta_c(G)$. The copy of s at u is a non-center node of $\mathcal{G}$ of degree at most $\delta_c(G) + \delta_n(H^*)$.

If $S = V(G)$ and $U = V(H^*)$ then the only possible non-center nodes belong to some H_s. A non-center node of H_s of degree at most $\delta_n(H_s)$ has degree at most $V(\mathcal{G}) - v(H) + \delta_n(H_s)$ in $\mathcal{G}$.

For δ_c the proofs are similar. If $S \neq V(G)$ then a center node w of G of degree $\delta_c(G)$ is not deleted in $\mathcal{G}$. If u is a center node of some H_s then the copy of t at u has degree at most $\delta_c(G) + \delta_c(H_s)$. On the other hand, if $S = V(G)$ then w may be deleted in the construction of $\mathcal{G}$, since w may in fact be the distinguished node v. In this case, if u is a node of H^* of minimal degree then the copy of w at u is a center node of degree at most $\delta_c(G) + \delta_c(H^*)$. □

In the case where both $v(G) = 2^a$ and $v(H) = 2^b$, Theorem 3.3 yields a center node set $\mathcal{S}$ equal to the entire set of nodes $V(\mathcal{G})$. If $v(\mathcal{G}) \neq 2^{a+b}$ then this is most likely not optimal (see Proposition 2.9), and yet this situation may be inevitable. For example, if $n = 2^m - 1$ then suppose that $n = xy - i$ and (x, y) satisfies the broadcast condition. Let $a = \lceil \log_2 x \rceil, b = \lceil \log_2 y \rceil$. Then $a + b = m$ and $xy - i = 2^m - 1$. The only solution is $x = 2^a, y = 2^b$ and $i = 1$, since if $x \leq 2^a - 1$ or $y \leq 2^b - 1$ then $xy \leq 2^m - 2$. Also, there is no solution with $i = 0$ and so the ordinary compounding method never generates an mbn with $2^m - 1$ nodes.

There is a slightly better method we can use in this special case, which reduces the center node sets obtained. Fix $a, b > 0$ with $a + b = m$. Proceed exactly as in the doubling procedure with G and H being mbn's with 2^a and 2^b nodes respectively, except that G^v may be replaced by any mbn G' on $2^a - 1$ nodes. Call the resulting graph $\mathcal{G}$. Let U be a center node set for H^* and W a center node set for G'.

Proposition 3.5. *In the notation of the above paragraph, the graph $\mathcal{G}$ is an mbn and $\mathcal{S} = V(\mathcal{G}) \setminus [(V(H^*) \setminus U) \cup V(G') \setminus W)]$ is a center node set for $\mathcal{G}$.*

Proof. Let $u \in V(\mathcal{G})$. If u is a node of some copy G_i then broadcast within G_i, taking a time steps. Then each node in G_i except possibly u is a center node and hence has an official message. Then broadcast horizontally in each copy of H, ensuring this is done officially in H^* with respect to U. This takes b time steps.

If u is a node of G' then first broadcast horizontally, which is possible since every node of G' also belongs to some copy of H. This takes b time steps. Every node receiving this message is a center node and so all nodes except possibly the originator now have the official message. Now broadcast vertically in the remaining a time steps, ensuring that this is done officially in G' with respect to W. □

Figure 6 illustrates this procedure. Here G H are each the hypercube Q_2, while G' and H^* are each the obn P_2 with 3 nodes. The white nodes are those which can be made non-center nodes according to Proposition 3.5.

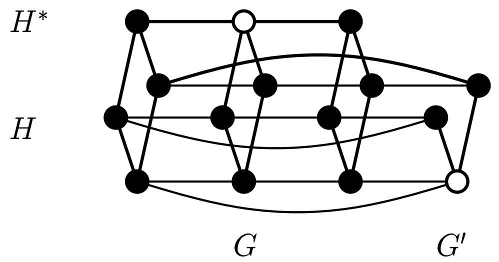

FIGURE 6. Example for Proposition 3.5

Proposition 3.6. *The following statements hold for the construction of Proposition 3.5.*

$$
\begin{aligned}
v(\mathcal{G}) &= 2^m - i \\
e(\mathcal{G}) &= m2^{m-1} + ie(G') + e(H^*) - (ia2^{a-1} + b2^{b-1}) \\
|\mathcal{S}| &= 2^m - i - [(v(H^*) - |U|) + (v(G') - |W|)] \\
\delta_n(\mathcal{G}) &\leq \min\{a + d_n(H^*), b + d_n(G'), m\} \\
\delta_c(\mathcal{G}) &\leq \min\{a + d_c(H^*), b + d_c(G'), m\}
\end{aligned}
$$

□

4. ITERATING THE COMPOUNDING METHODS

In order to generate mbn's with arbitrarily large number of nodes using the compound methods discussed above, it is necessary to iterate. Starting with a base table of mbn's and center node sets, one can generate new mbn's, each with a corresponding center node set, add these to the base table and repeat.

There will be gaps in the output for the ordinary compound method. For example, graphs whose order is prime or of the form $2^m - 1$ can not be generated. However, the doubling procedure generates graphs of all orders as long as the initial list contains the mbn's with 1 and 2 nodes. This last observation follows from the fact that the broadcast condition is always satisfied by $(2, n)$ and $(2, n, 1)$ for any positive integer n.

It is natural to suspect that the smaller center node sets given by Theorem 3.1 and Theorem 3.3 would lead to the mbn's generated at later stages having fewer edges than if the much larger center node sets given by Theorem 2.2 or Theorem 2.1 were used. Somewhat surprisingly, this is not the case, at least for the ordinary compounding method, as the following "associativity" result shows.

Say that two mbn's generated by the same iterated compounding method are equivalent if they have the same number of nodes and the same number of edges.

Proposition 4.1. *Any mbn eventually output from an initial list by iterating the ordinary compounding method with center node reduction as in Theorem 3.1 is equivalent to a compound of some mbn G into some mbn H, where G is in the initial list and H has been produced by the iterative method.*

PROOF. We claim first that if G, H and K be mbn's and let S and T be center node sets for G and H respectively then $(G_S[H])_{S\times T}[K]$ and $G_S[H_T[K]]$ are equivalent. This is a straightforward computation using the data in Proposition 3.2. The common vertex number is $v(G)v(H)v(K)$ and the common edge value $v(H)v(K)e(G) + |S|e(H)v(K) + |S||T|e(K)$. Note that if $(v(G), v(H))$ and $(v(G)v(H), v(K))$ satisfy the broadcast condition then it follows from the definition that $(v(H), v(K))$ and $(v(G), v(H)v(K))$ also do. A simple induction on the number of factors in a compound now proves the result. □

In the formula in Proposition 3.2 for $e(\mathcal{G})$ the size of a center node set for H does not appear. It follows from Proposition 4.1, by induction on the number of factors in a compound, that center node reduction of Theorem 3.1 does not improve the upper bounds on $B(n)$ generated by the iterated ordinary compounding method. However using center node reduction does result in the generation of many more non-isomorphic equivalent mbn's. This improves the chances of finding mbn's with "good" (for example, symmetric) broadcast protocols.

Empirical results suggest that the same phenomenon occurs with the doubling procedure, though we do not have a general proof of this fact. Proposition 4.1 can be generalized somewhat. We need only assume that both pairs (G, H) and (H, K) are combined as in the ordinary compound method. The general case seems much harder.

5. Computational results

In this section we discuss our implementation of the iterated compounding methods. The first subsection deals with the algorithms themselves, the second with the initial data used and the third with the empirical results obtained.

5.1. **The algorithms.** We have implemented two iterative algorithms as described in the last section. One uses ordinary compounding with the center node reduction as described in Theorem 3.1. The other combines the generalized compounding method as in Theorem 3.3, the special case center node reduction of Theorem 3.5 and the standard method of vertex deletion. The code for these programs is available from the first author on request.

It is infeasible to store and use all information about each known mbn. We use the data format of Table 1. Each row corresponds to a fixed mbn G with center node set $S(G)$. We allow the possibility of multiple rows for each $v(G)$ to take advantage of the structure of different mbn's and/or different center node sets.

Here $d_n(G)$ and $d_c(G)$ denote upper bounds for, respectively, $\delta_n(G)$ and $\delta_c(G)$, calculated with respect to $S(G)$. The column labeled 'Ref' gives a reference to where the given graph G (and its center node set) can be found. The bold entries in column $e(G)$ indicate when an mbn is an obn. In the same column, if the given graph $\mathcal{G}$ is generated by the doubling procedure, the relevant data $(v(G), v(H), i)$ or $(v(G), v(H))$ are listed. The G and H referred to also belong to the table.

Of course it is possible that better results might be obtained if more information were stored. For example, keeping the two lowest degrees of non-center nodes would improve some of the bounds in Proposition 3.4.

The first algorithm generates one new row for each ordered pair of input rows for which $(v(G), v(H))$ satisfies the broadcast criterion, using the update formula of Proposition 3.2.

For the second algorithm, rows are generated as follows. For each i with $0 \leq i \leq v(H) - 1$ the broadcast condition on $(v(G), v(H), i)$ is tested, and if satisfied a new row generated using the formulae of Proposition 3.4. Note that the expressions in those formulae are increasing functions of δ_c, δ_n and so the formulae give valid upper bounds if δ is replaced by d throughout. Note also that H^* can be any mbn with $v(H) - i$ nodes and so for H^* we use each mbn already available in the table with this property. We have already shown that there must already exist at least one such mbn in the table.

In addition, we generate rows by using Theorem 3.5 and the update formulae of Proposition 3.6 when and only when the resulting mbn has size $2^m - 1$.

Again it is infeasible to store and use all mbn's generated in each iteration. Since we are only concerned with the numbers of edges and center nodes it is reasonable to define a partial ordering $\preceq$ on the rows so that row $G_1 \preceq$ row G_2 if and only if $v(G_1) = v(G_2)$, $e(G_1) \leq e(G_2)$ and $c(G_1) \leq c(G_2)$. In this situation we say that G_1 *dominates* G_2. After a given iteration only the minimal rows with respect to this order are kept for the next iteration.

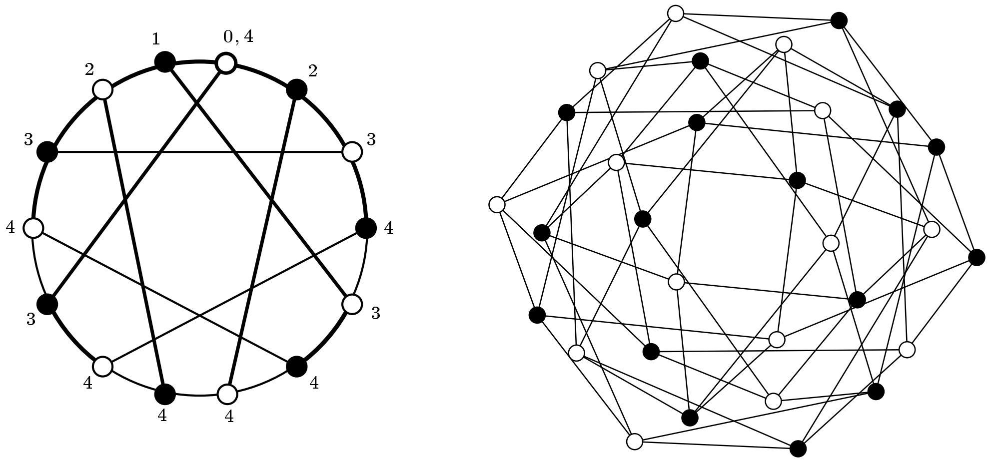

FIGURE 7. Bipartite vertex-transitive obn's with 14 and 30 nodes (see [6]).

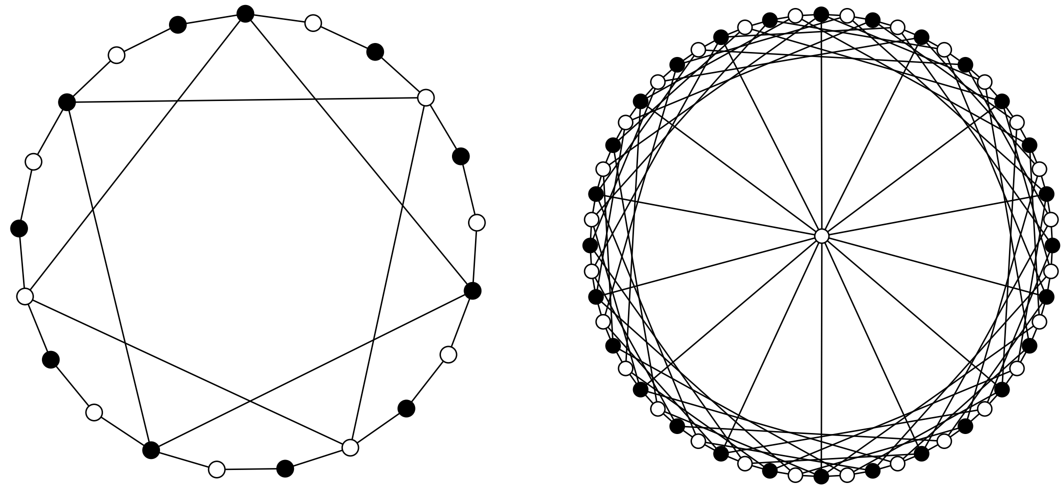

FIGURE 8. Mbn's with 21 and 57 nodes (see [3, 16])

For a preassigned positive integer M, the iteration in both algorithms stops when all rows with $v(G) \leq M$ show no change in two successive iterations. As a final step the second algorithm then tests each entry of the table to see whether vertex deletion improves the number of edges and if so generates a new row in this way.

5.2. **The initial data.** We used the mbn data listed in Table 1, with the addition of a few other mutually non-dominant mbn's that are not obn's (available by request). We now establish the validity of the entries in the initial table for which center node sets have not been given by other authors. As mentioned above the calculation of small center node sets is a nontrivial procedure, which has led to errors in the literature. Results such as Proposition 4.1 underline the importance of establishing good bounds on $cn(G)$ for the initial input.

The entries with 26–29 and 58–61 nodes are the obn's discovered by Saclé [20]. In that paper center node sets are given for the cases 26, 28 and 60. We treat the remaining cases here.

27: A vertex cover of size 15 with all nodes of degree at most 4 is shown in Figure 10.

29: The graph is displayed in Figure 10. By Proposition 2.8 all 16 nodes of degree 3 must lie in every cns. Estimating the size of a broadcast tree shows that every vertex of degree 4 must

$v(G)$	$e(G)$	$\|S(G)\|$	$d_n(G)$	$d_c(G)$	Ref
1	**0**	1		1	K_1
2	**1**	2		1	$Q_1 = K_2$
3	**2**	2	2	1	P_2
4	**4**	4		2	$Q_2 = (2,2)$
5	**5**	2	2	2	C_5
6	**6**	3	2	2	C_6
7	**8**	5	3	2	[23]
8	**12**	8		3	$Q_3 = (2,4)$
9	**10**	3	2	2	[23]
10	**12**	4	2	3	Fig. 2
11	**13**	5	2	2	[2]
12	**15**	6	2	3	(6,2)
13	**18**	7	3	2	[2]
14	**21**	7	3	3	Fig. 7
15	**24**	12	4	3	Fig. 1
16	**32**	16		4	$Q_4 = (2,8)$
17	**22**	5	2	3	[23]
18	**23**	6	2	3	(9,2)
19	**25**	7	2	2	[19]
20	**26**	8	2	3	[19]
21	**28**	11	2	3	Fig. 8
22	**31**	10	2	3	(11,2)
23	34	11	2	3	Fig. 9
24	36	11	3	3	[2]
25	40	15	3	3	[23]
26	**42**	15	3	3	[20]
27	**44**	15	3	3	Fig. 10
28	**48**	15	3	3	[20]
29	**52**	21	4	3	Fig. 10
30	**60**	15	4	4	Fig. 7
31	**65**	25	5	4	[23]
32	**80**	32		5	$Q_5 = (2,16)$
37	56	13	2	4	Fig. 9
39	59	16	2	4	Fig. 9
43	70	21	2	3	Fig. 9
49	94	30	3	4	[23]
57	126	28	4	4	Fig. 8
58	**121**	34	4	4	[20]
59	**124**	33	4	4	Fig. 11
60	**130**	30	4	4	Fig. 11
61	**136**	43	4	4	Fig. 12
62	**155**	31	5	5	[6]
63	**162**	54	6	5	[23]
64	**192**	64		6	$Q_6 = (2,32)$
$2^m - 2$	$\mathbf{(m-1)(2^{m-1}-1)}$	$2^{m-1} - 1$	m	m	[6]

TABLE 1. Standard input data.

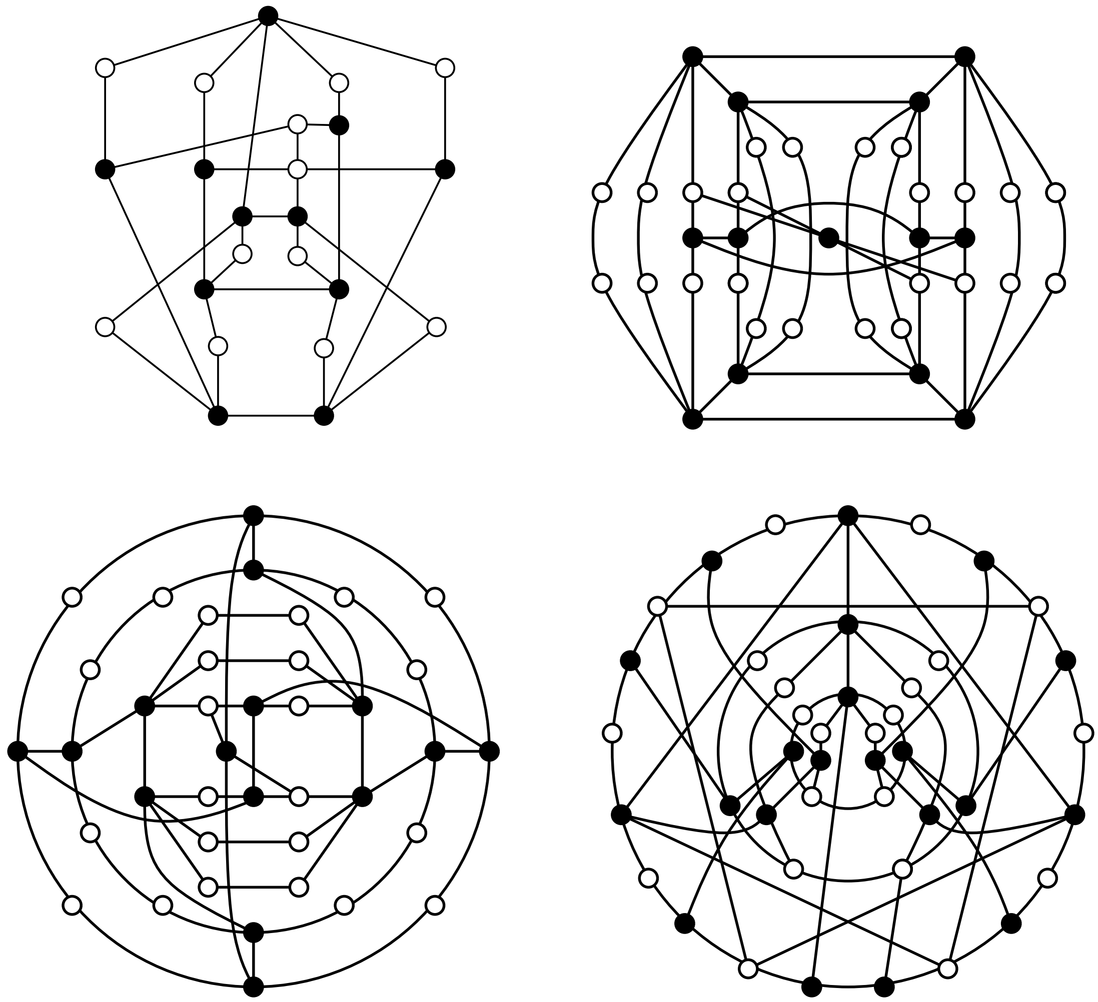

FIGURE 9. Mbn's with 23, 37, 39 and 43 nodes (see [19]).

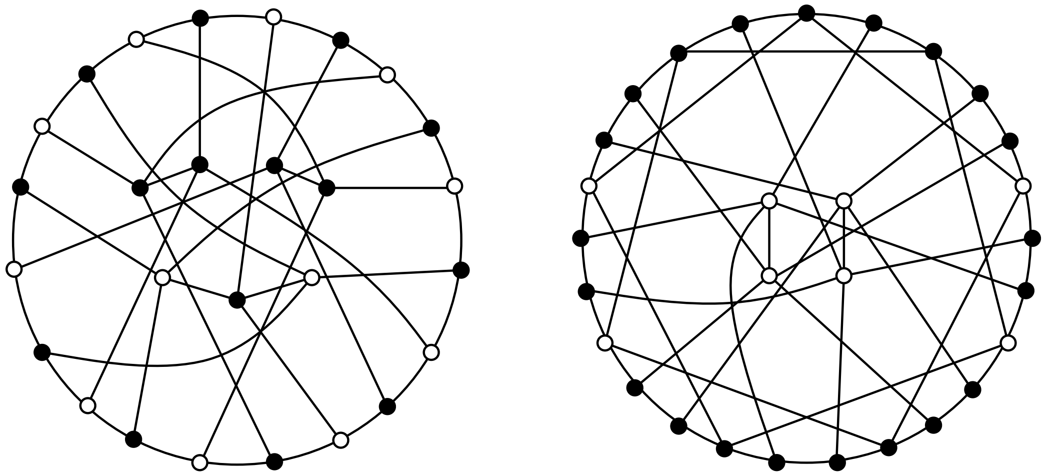

FIGURE 10. Obn's with 27 and 29 nodes (see [20]).

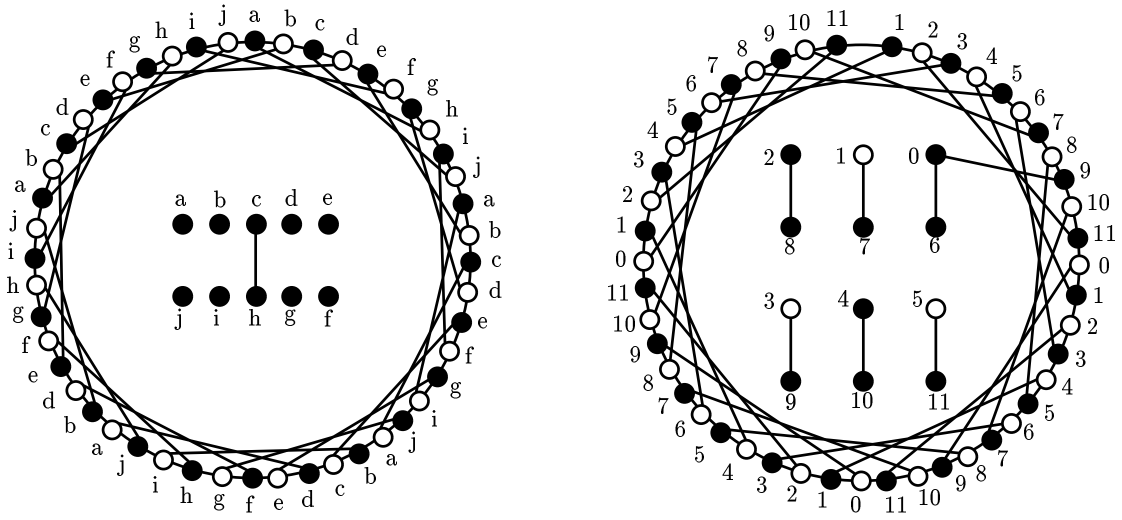

FIGURE 11. Obn's with 58 and 59 nodes (see [20]).

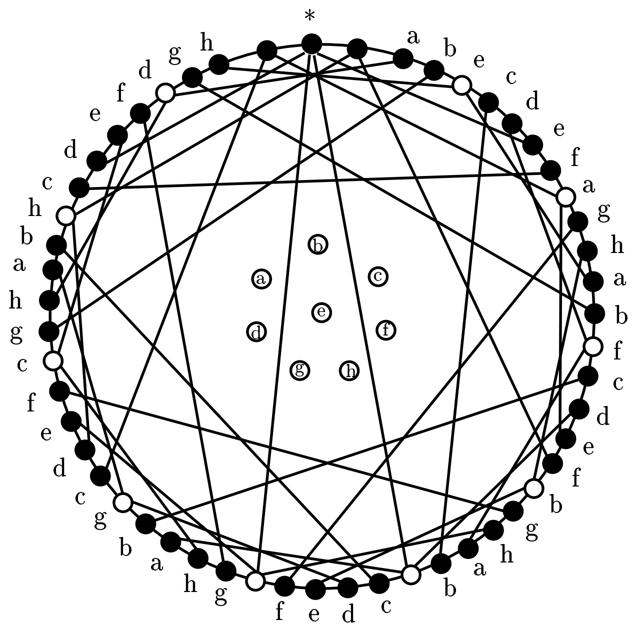

FIGURE 12. Obn with 61 nodes (see [20]).

send to another of degree 4 in every broadcast protocol. If a vertex of degree 4 has a neighbor of degree 4 which is not a center node then we obtain a contradiction since the maximum size of an official forest obtainable in such a situation is 24. Thus it is necessary to add to our candidate center node set a minimal vertex cover of the subgraph induced by the degree 4 nodes. Since these form a cycle of order 9 such a cover has size 5. We claim that the resulting set of size 21 is an ocns for G. If the originator is a non-center node of degree 4 then since all its neighbors are center nodes of degree at most 4 the result is clear. Now suppose that the originator has degree 5. Each such vertex v has one neighbor w of degree 5 and one of degree 3. In every broadcast protocol for v, either v informs a degree 3 node at time 1, or else v first informs w and then both v and w inform a degree 3 node at time 2. In either case there are enough degree 3 nodes idle to inform both v and w officially. All other nodes receive only the official message and so this case is completed.

58: The maximum degree is 5 and a solid 1-cover of size 34 is shown in Figure 11. For clarity not all edges have been drawn; each labeled node inside the main circle is connected to all nodes on the circle which have the same label.

59: The maximum degree is 5 and a solid 1-cover of size 33 is shown in Figure 11. The same convention is followed as in the previous case.

61: This is similar to case 29. The graph is displayed in Figure 12, with the same convention as in the previous two cases. By Proposition 2.8 all nodes of degree 4 must lie in every cns. Adding the vertex of degree 6 labeled '*', we obtain a set of size 43 which the following observations show is a center node set.

If the originator v has degree 5 then estimating the size of a broadcast tree shows that its first message must go to a w of node of degree 6 and its second to a node x of degree 4. Then x is idle at time at most 6 and so can inform v officially. If w is the node labeled '*' then there is nothing left to prove. Otherwise w sends to a degree 4 node $y \neq x$ at time 2, and this is then idle at time at most 6 and can inform w officially.

If the originator is a non-center node of degree 6 then it must send to a degree 4 node at time at most 2 and the result follows as above.

The obn's with $2^m - 2$ nodes presented in [6] are regular bipartite and of degree $m-1$, with either half of the bipartition yielding a vertex cover. The obn with 21 nodes and 28 edges of [16] has maximum degree 4 and so a solid 1-cover for this graph is shown in Figure 8. The mbn with 57 nodes and 126 edges from [3] has a solid 1-cover with all nodes of degree at most 5 and is also shown in Figure 8.

The mbn's discovered by Ridwan [19] are shown in Figure 9 along with the center node sets given in that paper.

5.3. **Empirical results.** Programs implementing each algorithm described above were run subject to the added heuristic (for memory reasons) that after each iteration, at most 10 mutually nondominant rows were used in the next iteration, for each fixed number n of nodes. The larger number of combinations which must be considered in the doubling procedure was reflected in the fact the running time for the first algorithm was only a few minutes as opposed to several hours for the second. Computations in this paper were carried out using the University of Auckland's SGI Power Challenge GR computer (with 16 R10000 processors).

In the range $n \leq 16384$, the first (ordinary compounding) algorithm generated upper bounds for $B(n)$ for 1618 of the values of n. For these values of n the second algorithm lowered those bounds 75.3% of the time, with a mean improvement of 9.42%.

In the same range the center node reduction techniques in this paper achieved a mean reduction of around 50% in the upper bounds for $cn(G)$ formerly produced by the compounding methods.

In [20] lower bounds for $B(n)$ were presented, where n has the form $2^m - 3, 2^m - 4, 2^m - 5$ or $2^m - 6$. For these values of n less than 16384 the best upper bounds of this paper exceed the lower bounds by an average of less than 6%. Most notably it follows from Saclé's lower bound and our upper bound that $314 \leq B(122) \leq 315$.

The best upper bounds for $B(n)$ generated by various methods, for certain ranges of n, are displayed in Table 2. These sample values of n were established by Gargano and Vaccaro [10]. The second column (*WV* or Ref) indicates the previous known bound before Bermond et al. [2]. The default reference *WV* is [23]. The *BFP* and *DVWZ* columns indicate the best known bounds obtained in this current paper, by our first (ordinary compounding) and second algorithms respectively. Only the cases $n = 513, 896, 1008, 16128$ were presented in [2], and the corresponding bounds on $B(n)$ were the same as in this paper. We note that the justification for the *BFP* entry given in

[2] for the $n = 1008$ row was incorrect, being based on the assumption that Labahn's obn G with 63 nodes has a solid 1-cover of size 36. In fact $cn(G) = 54$ by Proposition 2.8 and this gives a value of 4320 for the particular compound used. However, one may instead take G to be the regular obn with 126 nodes and H to be the hypercube Q_3. This yields $B(1008) \leq 3780$, the entry shown in Table 2.

Many authors have published tables of the best known upper bounds for $B(n)$ for $1 \leq n \leq 64$. Table 3 shows the best known upper bounds for $B(n)$ for $17 \leq n \leq 127$. The (exact) bounds for $n \leq 16$ (which are listed in Table 1) are credited to Farley et al. [7, 8].

Each entry in the column labeled 'Compound' gives a construction that matches the best known bound. In every case except for $n = 114$, the G and H in the entry $(v(G), v(H), i)$ also occur in Table 3. In this special case the algorithm uses the graph G with 57 nodes from Table 1, which does not yield the best bound on $B(57)$. This shows the usefulness of keeping several mutually nondominant mbn's, rather than taking just the one with the fewest edges, and underlines the importance of finding small center node sets.

New bounds for $B(n)$ are indicated in each row of Table 3 where we omit a reference. Again bold entries in the table denote optimal values. We note that the entry 111 for $n = 55$ was claimed in [23] but cannot be obtained via the doubling procedure with the data listed in that paper.

6. Comments

The most difficult case for any of the methods above, as far as generating an upper bound for $B(n)$ is concerned, is the case where $n = 2^m - 1$. Despite our refinements the results in this paper are rather far from the conjectured values of $B(n)$. It seems necessary to attack this case directly.

The widely believed conjecture that $B(n) \leq B(n+1)$ if $n \neq 2^m$ is still unproven.

As mentioned after Proposition 4.1, it is unknown whether the iterated doubling procedure can give better upper bounds on $B(n)$ when center node reduction is used.

Acknowledgments. We thank Art Liestman for providing his current table of best known values of $B(n)$, $n \leq 64$. The second author was partially supported by the National Science Foundation under Grant DDM 90-57066. The third author was supported by a NZST Postdoctoral Fellowship.

n	VW or Ref	BFP	$DVWZ$
513	1026	948	934
514	1028		935
515	1030		936
516	1032	999	937
517	1034		938
520	1040	976	941
521	1042		942
528	1056	1020	959
529	1058	1156	960
544	1088	1104	1013
545	1090		1014
576	1152	1140	1128
577	1410		1129
640	1536	1472	1348
641	1729		1349
768	2112	2032	1932
769	2690		1932
896	2880	2688	2678
897	3264		2676
960	3840	3040	3040
961	4040	3640	3479
992	4512	3472	3472
993	4616		3900
1008	4832	3780	3780
1009	4889		4197
1016	4984	4064	4064
1017	5017		4466
1020	5056 [10]	4335	4335
1021	5078		4626
1022	4599 [6, 14]	—	—
1023	5106		5082

n	VW or Ref	BFP	$DVWZ$
8193	20482		17083
8194	20484		17084
8195	20486		17085
8196	20488		17086
8197	20490		17087
8200	20496		17090
8201	20498		17091
8208	20512	17856	17098
8209	20514		17099
8224	20544		17114
8225	20546		17115
8256	20608	20016	17146
8257	20610		17147
8320	20736	19200	17406
8321	20738		17407
8448	20992	20144	17968
8449	20994		17969
8704	21504	22784	19260
8705	21506		19261
9216	22528	22464	21176
9217	29698		21177
10240	32768	31744	24541
10241	37889		24542
12288	45057	43008	35637
12289	59394		35637
14336	62464	57344	49748
14337	68610		49746
15360	81920 [10]	64000	61584
15361	89093 [10]		61579
15872	91648 [10]	71424	71424
15873	99334 [10]		71469
16128	100864 [10]	76608	76608
16129	108807 [10]		83464
16256	109696 [10]	81280	81280
16257	111883		87939
16320	113152 [10]	85680	85680
16321	113356		92169
16352	113952	89936	89936
16353	114060		96294
16368	114336	94116	94116
16369	114397		100379
16376	114520	98256	98256
16377	114557		104452
16380	114608	102375	102375
16381	114648 [7]		106548
16382	106483 [6, 14]	—	—
16383	114674 [7]		114626

TABLE 2. Comparison of recent methods for bounding $B(n)$, $n = 2^9 \sim 2^{10}$ and $n = 2^{13} \sim 2^{14}$.

n	Bound	Ref	Compound	n	Bound	Ref	Compound	n	Bound	Ref	Compound
17	**22**	[17]		54	103		(27,2)	91	179		(23,4,1)
18	**23**	[3]	(9,2)	55	111		(28,2,1)	92	180		(23,4)
19	**25**	[3]		56	111	[20]	(28,2)	93	188		(24,4,3)
20	**26**	[16]		57	123		(58,1,1)	94	188		(24,4,2)
21	**28**	[16]		58	**121**	[20]		95	188		(24,4,1)
22	**31**	[16]	(11,2)	59	**124**	[20]		96	188		(24,4)
23	34	[16]		60	**130**	[20]		97	203		(14,7,1)
24	36	[3]		61	**136**	[20]		98	203		(14,7)
25	40	[3]		62	**155**	[6, 14]		99	220		(25,4,1)
26	**42**	[20]		63	**162**	[15]		100	220		(25,4)
27	**44**	[20]		64	**192**	[8]	(2,32)	101	228		(13,8,3)
28	**48**	[20]		65	101		(5,13)	102	228		(13,8,2)
29	**52**	[20]		66	105		(6,11)	103	228		(13,8,1)
30	**60**	[3]		67	107		(17,4,1)	104	228		(13,8)
31	**65**	[3]		68	108		(17,4)	105	236		(27,4,3)
32	**80**	[8]	(2,16)	69	111		(5,14,1)	106	236		(27,4,2)
33	48	[23]	(3,11)	70	112		(5,14)	107	236		(27,4,1)
34	49	[23]	(17,2)	71	115		(9,8,1)	108	236		(27,4)
35	51	[2]	(5,7)	72	116		(9,8)	109	252		(14,8,3)
36	52	[2]	(9,4)	73	121		(5,15,2)	110	252		(14,8,2)
37	56	[19]		74	122		(5,15,1)	111	252		(14,8,1)
38	57	[3]	(19,2)	75	123		(5,15)	112	252		(14,8)
39	59	[19]		76	128		(19,4)	113	281		(2,59,5)
40	60	[3]	(20,2)	77	131		(6,13,1)	114	280		(57,2)
41	65	[2]	(6,7,1)	78	132		(6,13)	115	278		(58,2,1)
42	66	[2]	(6,7)	79	135		(20,4,1)	116	276		(58,2)
43	70	[19]		80	136		(20,4)	117	283		(59,2,1)
44	72	[2]	(11,4)	81	142		(3,27)	118	281		(59,2)
45	78		(23,2,1)	82	145		(6,14,2)	119	292		(60,2,1)
46	79		(23,2)	83	146		(6,14,1)	120	290	[20]	(60,2)
47	83	[2]	(24,2,1)	84	147		(6,14)	121	317		(61,2,1)
48	83	[2]	(24,2)	85	157		(6,15,5)	122	315		(61,2)
49	94	[23]		86	158		(6,15,4)	123	346		(62,2,1)
50	95	[2]	(25,2)	87	159		(6,15,3)	124	341		(62,2)
51	99		(26,2,1)	88	160		(6,15,2)	125	379		(2,63,1)
52	99	[20]	(26,2)	89	161		(6,15,1)	126	**378**	[6, 14]	
53	103		(27,2,1)	90	162		(6,15)	127	417		(2,64,1)

TABLE 3. Best known upper bounds on $B(n)$, $17 \leq n \leq 127$.

References

1. M. Barborak, M. Malek, and A. Dahbura, *The consensus problem for fault-tolerant computing*, ACM Comput. Surv. **25** (1993), 171–220.
2. J.-C. Bermond, P. Fraigniaud, and J. G. Peters, *Antepenultimate broadcasting*, Networks **26** (1995), 125–137.
3. J.-C. Bermond, P. Hell, A. L. Liestman, and J. G. Peters, *Sparse broadcast graphs*, Discrete Appl. Math. **36** (1992), 97–130.
4. S. C. Chau and A. L. Liestman, *Constructing minimal broadcast networks*, J. Combin. Inform. Syst. Sci. **10** (1985), 110–122.
5. X. Chen, *An upper bound for the broadcast function* $B(n)$, Chinese J. Computers **13** (1990), 605–611.
6. M. J. Dinneen, M. R. Fellows, and V. Faber, *Algebraic constructions of efficient broadcast networks*, in Applied Algebra, Algebraic Algorithms and Error-Correcting Codes 9, Lecture Notes in Computer Science, vol. 539, 1991, pp. 152–158.
7. A. M. Farley, *Minimal broadcast networks*, Networks **9** (1979), 313–332.
8. A. M. Farley, S. T. Hedetniemi, A. Proskurowski, and S. Mitchell, *Minimum broadcast graphs*, Discrete Math. **25** (1979), 189–193.
9. P. Fraigniaud and E. Lazard, *Methods and problems of communication in usual networks*, Discrete Appl. Math. **53** (1994), 79–133.
10. L. Gargano and U. Vaccaro, *On the construction of minimal broadcast networks*, Networks **19** (1989), 673–689.
11. M. Grigni and D. Peleg, *Tight bounds on minimum broadcast networks*, SIAM J. Discrete Math. **4** (1991), 207–222.
12. S. M. Hedetniemi, S. T. Hedetniemi, and A. L. Liestman, *A survey of gossiping and broadcasting in communication networks*, Networks **18** (1988), 319–349.
13. J. Hromkovič, R. Klasing, B. Monien, and R. Peine, *Dissemination of information in interconnection networks (broadcasting and gossiping)*, in Combinatorial Network Theory, Applied Optimization, vol. 1, Kluwer, 1996, pp. 125–212.
14. L. H. Khachatrian and O. S. Harutounian, *Construction of new classes of minimal broadcast networks*, Conference on Coding Theory, 1990.
15. R. Labahn, *A minimum broadcast graph on 63 vertices*, Discrete Appl. Math. **53** (1994), 247–250.
16. M. Maheo and J.-F. Saclé, *Some minimum broadcast graphs*, Discrete Appl. Math. **53** (1994), 275–285.
17. S. Mitchell and S. Hedetniemi, *A census of minimum broadcast graphs*, J. Combin. Inform. System Sci. **5** (1980), 141–151.
18. A. Pelc, *Fault-tolerant broadcasting and gossiping in communication networks*, Networks **28** (1996), 143–156.
19. Ridwan, *Construction of minimal broadcast networks*, Master's thesis, Pennsylvania State University, August 1995.
20. J.-F. Saclé, *Lower bounds for the size in four families of minimum broadcast graphs*, Discrete Math. **150** (1996), 359–369.
21. P. J. Slater, E. Cockayne, and S. T. Hedetniemi, *Information dissemination in trees*, SIAM J. Comput. **10** (1981), 692–701.
22. J. A. Ventura and X. Weng, *A new method for constructing minimal broadcast networks*, Networks **23** (1993), 481–497.
23. M. X. Weng and J. A. Ventura, *A doubling procedure for constructing minimal broadcast networks*, Telecomm. Syst. **3** (1995), 259–293.

Department of Computer Science, University of Auckland, Private Bag 92019 Auckland, New Zealand
E-mail address: `mjd@cs.auckland.ac.nz`

Department of Industrial and Manufacturing Engineering, The Pennsylvania State University, University Park, Pennsylvania 16802, USA
E-mail address: `jav1@psu.edu`

Department of Mathematics, University of Auckland, Private Bag 92019 Auckland, New Zealand
E-mail address: `wilson@math.auckland.ac.nz`

Department of Engineering Science, University of Auckland, Private Bag 92019 Auckland, New Zealand
E-mail address: `g.zakeri@auckland.ac.nz`